arXiv:1403.7703v1 [math.NT] 30 Mar 2014

General systems of linear forms: equidistribution and true complexity

Hamed Hatami * Pooya Hatami † Shachar Lovett ‡

April 1, 2014

Abstract

The densities of small linear structures (such as arithmetic progressions) in subsets of Abelian groups can be expressed as certain analytic averages involving linear forms. Higher-order Fourier analysis examines such averages by approximating the indicator function of a subset by a function of bounded number of polynomials. Then, to approximate the average, it suffices to know the joint distribution of the polynomials applied to the linear forms. We prove a near-equidistribution theorem that describes these distributions for the group $\mathbb{F}_p^n$ when p is a fixed prime. This fundamental fact is equivalent to a strong near-orthogonality statement regarding the higher-order characters, and was previously known only under various extra assumptions about the linear forms.

As an application of our near-equidistribution theorem, we settle a conjecture of Gowers and Wolf on the true complexity of systems of linear forms for the group $\mathbb{F}_p^n$.

1 Introduction

Gowers' seminal work in combinatorial number theory [Gow01] initiated an extension of the classical Fourier analysis, called *higher-order Fourier analysis* of Abelian groups. Higher-order Fourier analysis has been very successful in dealing with problems regarding the densities of small linear structures (e.g. arithmetic progressions) in subsets of Abelian groups. It is possible to express such densities as certain analytic averages. For example, the density of the three term arithmetic progressions in a subset A of an Abelian group G can be expressed as $\mathbf{E}_{x,y\in G}\left[\mathbf{1}_A(x)\mathbf{1}_A(x+y)\mathbf{1}_A(x+2y)\right]$. More generally, one is often interested in analyzing

$$\mathop{\mathbf{E}}_{x_1,\ldots,x_k\in G}\left[\mathbf{1}_A(L_1(x_1,\ldots,x_k))\cdots\mathbf{1}_A(L_m(x_1,\ldots,x_k))\right], \tag{1}$$

where each L_i is a linear form on k variables. Averages of this type are of interest in computer science, additive combinatorics, and analytic number theory.

In this paper we are only interested in the group $\mathbb{F}^n$ where $\mathbb{F}=\mathbb{F}_p$ for a fixed prime p and n is large. In the classical Fourier analysis of $\mathbb{F}^n$, a function is expressed as a linear combination of the characters of $\mathbb{F}^n$. Note that the characters of $\mathbb{F}^n$ are exponentials of linear polynomials:

*McGill University. hatami@cs.mcgill.ca. Supported by an NSERC, and an FQRNT grant.
†University of Chicago. pooya@cs.uchicago.edu.
‡UC San Diego. slovett@ucsd.edu. Supported by NSF CAREER award 1350481.

for $\alpha \in \mathbb{F}^n$, the corresponding character is defined as $\chi_\alpha(x) = \mathsf{e}\left(\sum_{i=1}^n \alpha_i x_i\right)$, where $\mathsf{e}(a) := e^{\frac{2\pi i}{p}a}$ for $a \in \mathbb{F}$. In higher-order Fourier analysis, the linear polynomials are replaced by higher degree polynomials, and one would like to approximate a function $f : \mathbb{F}^n \to \mathbb{C}$ by a linear combination of the functions $\mathsf{e}(P)$, where each P is a polynomial of a certain degree. The existence of such approximations is a consequence of the the so-called "inverse theorems" for Gowers norms which are established in a sequence of papers by Bergelson, Green, Samorodnitsky, Szegedy, Tao, and Ziegler [TZ10, TZ12, Sze12, GTZ12, GT10a, BTZ10, Sam07].

Higher-order Fourier expansions are extremely useful in studying averages that are defined through linear structures. To analyze the average in Equation (1), one approximates $1_A \approx \Gamma(P_1, \ldots, P_C)$ where $\Gamma : \mathbb{F}^C \to \mathbb{C}$ is a function that is applied to a constant number of low degree polynomials $P_1, \ldots, P_C : \mathbb{F}^n \to \mathbb{F}$. Then applying the classical Fourier transform to Γ yields the higher-order Fourier expansion

$$\mathbf{1}_A \approx \sum_{\alpha \in \mathbb{F}^C} \widehat{\Gamma}(\alpha)\, \mathsf{e}\left(\sum_{i=1}^n \alpha_i P_i\right),$$

where the coefficients $\widehat{\Gamma}(\alpha)$ are complex numbers.

Near-Orthogonality and Equidistribution. One of the important and useful properties of the classical Fourier characters is that they form an orthonormal basis. For higher-order Fourier expansions to be useful, one needs a similar orthogonality for the higher-order characters $\mathsf{e}\left(\sum_{i=1}^n \alpha_i P_i\right)$ appearing in the expansion, or at least an approximation of it. This approximate orthogonality is established by Green and Tao [GT09] and Kaufman and Lovett [KL08], and is in fact equivalent to a near-equidistribution statement: the polynomials in the approximation $1_A \approx \Gamma(P_1, \ldots, P_C)$ can be chosen in such a way that the distribution of $(P_1(x), \ldots, P_C(x))$ is close to the uniform distribution on $\mathbb{F}^C$ when x is chosen uniformly at random from $\mathbb{F}^n$.

However, this is not completely satisfactory, as to study the averages of the form Equation (1), one needs to understand the distribution of the more sophisticated random variable

$$\begin{pmatrix} P_1(L_1(X)) & P_2(L_1(X)) & \ldots & P_C(L_1(X)) \\ P_1(L_2(X)) & P_2(L_2(X)) & \ldots & P_C(L_2(X)) \\ \vdots & & & \vdots \\ P_1(L_m(X)) & P_2(L_m(X)) & \ldots & P_C(L_m(X)) \end{pmatrix},$$

where $X = (x_1, \ldots, x_k)$ is the uniform random variable taking values in $(\mathbb{F}^n)^k$. Since polynomials of a given degree satisfy various linear identities (e.g. every degree one polynomial P satisfies $P(x+y+z) = P(x+y) + P(x+z) - P(x)$), it is no longer possible to choose the polynomials in a way that this random matrix is almost uniformly distributed on $\mathbb{F}^{C \times m}$. Therefore, in this case one would like to obtain an almost uniform distribution on the points of $\mathbb{F}^{C \times m}$ that are consistent with these linear identities. Note that while [GT09, KL08] only say that the entries in each row of this matrix are nearly independent, such a stronger near-equidistribution would in particular imply that the columns of this matrix are nearly independent.

Hatami and Lovett [HL11] established this strong near-equidistribution in the case where the characteristic of the field $\mathbb{F}$ is greater than the degree of the involved polynomials. Bhattacharyya, *et al.* [BFH+13] extended the result of [HL11] to the general characteristic case, but under the

extra assumption that the system of linear forms is affine, i.e. there is a variable that appears with coefficient 1 in all the linear forms. Finally, in the present paper, in Theorem 3.7 we prove the near-equidistribution statement without any extra assumptions on the linear forms.

A conjecture of Gowers and Wolf. In dealing with the averages of the form Equation (1) a question arises naturally: Given such an average, what is the smallest k such that there is an approximation of $\mathbf{1}_A$ with a linear combination of a few higher-order characters of degree at most k that affects the average only negligibly? This question was asked and studied by Gowers and Wolf [GW10] who conjectured a simple characterization for this value, and verified it for the case of large $|\mathbb{F}|$ in [GW11]. As an application of our near-orthogonality result, we settle the Gowers-Wolf conjecture in full generality on $\mathbb{F}^n$. In the setting of functions on $\mathbb{Z}_N$, Green and Tao [GT10a] established similar results and characterizations.

Homogeneous non-classical polynomials. The main difficulty in dealing with fields of low characteristic is that in the higher-order Fourier expansions, instead of the exponentials of classical polynomials, one has to work with exponentials of a generalization of them which are referred to as "non-classical" polynomials. Recall that a classical polynomial is homogeneous if all of its monomials are of the same degree. A useful property of a homogeneous classical polynomial $P(x)$ of degree d is that $P(cx) = c^d P(x)$, for every $c \in \mathbb{F}$. We use this property to extend the definition of homogeneity to non-classical polynomials. An ingredient of the proof of our near-equidistribution result is a statement about non-classical polynomials which we believe is of independent interest. In Theorem 3.4 we show that homogeneous multivariate (non-classical) polynomials span the space of multivariate (non-classical) polynomials. We later use this to prove our near-equidistribution results for homogeneous polynomials.

2 Notation and Preliminaries

Fix a prime field $\mathbb{F} = \mathbb{F}_p$ for a prime $p \geqslant 2$. Throughout the paper, we fix $\zeta \in \mathbb{F}^*$ a generator of $\mathbb{F}^*$. Define $|\cdot|$ to be the standard map from $\mathbb{F}$ to $\{0, 1, \ldots, p-1\} \subset \mathbb{Z}$. Let $\mathbb{D}$ denote the complex unit disk $\{z \in \mathbb{C} : |z| \leqslant 1\}$.

For integers a, b, we let $[a]$ denote the set $\{1, 2, \ldots, a\}$ and $[a, b]$ denote the set $\{a, a+1, \ldots, b\}$. For real numbers, $\alpha, \sigma, \varepsilon$, we use the shorthand $\sigma = \alpha \pm \varepsilon$ to denote $\alpha - \varepsilon \leqslant \sigma \leqslant \alpha + \varepsilon$. The power set of a set S is denoted by $\mathcal{P}(S)$. The zero element in $\mathbb{F}^n$ is denoted by $\underline{0}$. We will denote by lower case letters, e.g. x, y, elements of $\mathbb{F}^n$. We use capital letters, e.g. $X = (x_1, \ldots, x_k) \in (\mathbb{F}^n)^k$, to denote tuples of variables.

Definition 2.1. *A linear form on k variables is a vector $L = (\ell_1, \ldots, \ell_k) \in \mathbb{F}^k$ and it maps $X = (x_1, \ldots, x_k) \in (\mathbb{F}^n)^k$ to $L(X) = \sum_{i=1}^k \ell_i x_i \in \mathbb{F}^n$.*

For a linear form $L = (\ell_1, \ldots, \ell_k) \in \mathbb{F}^k$ we define $|L| \stackrel{\text{def}}{=} \sum_{i=1}^k |\ell_i|$.

2.1 Higher-order Fourier Analysis

We need to recall some definitions and results about higher-order Fourier analysis. Most of the material in this section is directly quoted from the full version of [BFH+13].

Definition 2.2 (Multiplicative Derivative)**.** *Given a function $f : \mathbb{F}^n \to \mathbb{C}$ and an element $h \in \mathbb{F}^n$, define the* multiplicative derivative *in direction h of f to be the function $\Delta_h f : \mathbb{F}^n \to \mathbb{C}$ satisfying $\Delta_h f(x) = f(x+h)\overline{f(x)}$ for all $x \in \mathbb{F}^n$.*

The *Gowers norm* of order d for a function $f : \mathbb{F}^n \to \mathbb{C}$ is the expected multiplicative derivative of f in d random directions at a random point.

Definition 2.3 (Gowers norm)**.** *Given a function $f : \mathbb{F}^n \to \mathbb{C}$ and an integer $d \geqslant 1$, the* Gowers norm of order d *for f is given by*

$$\|f\|_{U^d} = \left| \mathop{\mathbf{E}}_{y_1,\ldots,y_d,x \in \mathbb{F}^n} [(\Delta_{y_1}\Delta_{y_2}\cdots\Delta_{y_d} f)(x)] \right|^{1/2^d}.$$

Note that as $\|f\|_{U^1} = |\mathbf{E}[f]|$ the Gowers norm of order 1 is only a semi-norm. However for $d > 1$, it is not difficult to show that $\|\cdot\|_{U^d}$ is indeed a norm.

If $f = e^{2\pi i P/p}$ where $P : \mathbb{F}^n \to \mathbb{F}$ is a polynomial of degree $< d$, then $\|f\|_{U^d} = 1$. If $d < p$ and $\|f\|_\infty \leqslant 1$, then in fact, the converse holds, meaning that any function $f : \mathbb{F}^n \to \mathbb{C}$ satisfying $\|f\|_\infty \leqslant 1$ and $\|f\|_{U^d} = 1$ is of this form. But when $d \geqslant p$, the converse is no longer true. In order to characterize functions $f : \mathbb{F}^n \to \mathbb{C}$ with $\|f\|_\infty \leqslant 1$ and $\|f\|_{U^d} = 1$, one needs to define the notion of *non-classical polynomials*.

Non-classical polynomials might not be necessarily $\mathbb{F}$-valued. We need to introduce some notation. Let $\mathbb{T}$ denote the circle group $\mathbb{R}/\mathbb{Z}$. This is an Abelian group with group operation denoted $+$. For an integer $k \geqslant 0$, consider the subgroup $\frac{1}{p^k}\mathbb{Z}/\mathbb{Z} \subseteq \mathbb{T}$. Let $\mathsf{e} : \mathbb{T} \to \mathbb{C}$ denote the character $\mathsf{e}(x) = e^{2\pi i x}$.

Definition 2.4 (Additive Derivative)**.** *Given a function[1] $P : \mathbb{F}^n \to \mathbb{T}$ and an element $h \in \mathbb{F}^n$, define the* additive derivative *in direction h of f to be the function $D_h P : \mathbb{F}^n \to \mathbb{T}$ satisfying $D_h P(x) = P(x+h) - P(x)$ for all $x \in \mathbb{F}^n$.*

Definition 2.5 (Non-classical polynomials)**.** *For an integer $d \geqslant 0$, a function $P : \mathbb{F}^n \to \mathbb{T}$ is said to be a* non-classical polynomial of degree $\leqslant d$ *(or simply a* polynomial of degree $\leqslant d$*) if for all $y_1, \ldots, y_{d+1}, x \in \mathbb{F}^n$, it holds that*

$$(D_{y_1} \cdots D_{y_{d+1}} P)(x) = 0. \tag{2}$$

The degree *of P is the smallest d for which the above holds. A function $P : \mathbb{F}^n \to \mathbb{T}$ is said to be a* classical polynomial of degree $\leqslant d$ *if it is a non-classical polynomial of degree $\leqslant d$ whose image is contained in $\frac{1}{p}\mathbb{Z}/\mathbb{Z}$.*

It is a direct consequence of the definition that a function $f : \mathbb{F}^n \to \mathbb{C}$ with $\|f\|_\infty \leqslant 1$ satisfies $\|f\|_{U^{d+1}} = 1$ if and only if $f = \mathsf{e}(P)$ for a (non-classical) polynomial $P : \mathbb{F}^n \to \mathbb{T}$ of degree $\leqslant d$. We denote by $\mathrm{Poly}(\mathbb{F}^n \to \mathbb{T})$ and $\mathrm{Poly}_{\leqslant d}(\mathbb{F}^n \to \mathbb{T})$, respectively, the set of all non-classical polynomials, and the ones of degree at most d.

The following lemma of Tao and Ziegler [TZ12] shows that a classical polynomial P of degree d must always be of the form $x \mapsto \frac{\lfloor Q(x) \rfloor}{p}$, where $Q : \mathbb{F}^n \to \mathbb{F}$ is a polynomial (in the usual sense) of degree d, and $\lfloor \cdot \rfloor$ is the standard map from $\mathbb{F}$ to $\{0, 1, \ldots, p-1\}$. This lemma also characterizes the structure of non-classical polynomials.

[1] We try to adhere to the following convention: upper-case letters (e.g. F and P) to denote functions mapping from $\mathbb{F}^n$ to $\mathbb{T}$ or to $\mathbb{F}$, lower-case letters (e.g. f and g) to denote functions mapping from $\mathbb{F}^n$ to $\mathbb{C}$, and upper-case Greek letters (e.g. Γ and Σ) to denote functions mapping $\mathbb{T}^C$ to $\mathbb{T}$.

Lemma 2.6 (Lemma 1.7 in [TZ12]). *A function $P : \mathbb{F}^n \to \mathbb{T}$ is a polynomial of degree $\leqslant d$ if and only if P can be represented as*

$$P(x_1, \dots, x_n) = \alpha + \sum_{\substack{0 \leqslant d_1, \dots, d_n < p; k \geqslant 0: \\ 0 < \sum_i d_i \leqslant d - k(p-1)}} \frac{c_{d_1,\dots,d_n,k} |x_1|^{d_1} \cdots |x_n|^{d_n}}{p^{k+1}} \mod 1,$$

for a unique choice of $c_{d_1,\dots,d_n,k} \in \{0, 1, \dots, p-1\}$ and $\alpha \in \mathbb{T}$. The element α is called the shift *of P, and the largest integer k such that there exist $d_1, \dots, d_n$ for which $c_{d_1,\dots,d_n,k} \neq 0$ is called the* depth *of P. A depth-k polynomial P takes values in a coset of the subgroup $\mathbb{U}_{k+1} \stackrel{\text{def}}{=} \frac{1}{p^{k+1}}\mathbb{Z}/\mathbb{Z}$. Classical polynomials correspond to polynomials with 0 shift and 0 depth.*

Note that Lemma 2.6 immediately implies the following important observation[2]:

Remark 2.7. *If $Q : \mathbb{F}^n \to \mathbb{T}$ is a polynomial of degree d and depth k, then pQ is a polynomial of degree $\max(d - p + 1, 0)$ and depth $k - 1$. In other words, if Q is classical, then pQ vanishes, and otherwise, its degree decreases by $p - 1$ and its depth by 1. Also, if $\lambda \in [1, p-1]$ is an integer, then $\deg(\lambda Q) = d$ and $\operatorname{depth}(\lambda Q) = k$.*

For convenience of exposition, we will assume throughout this paper that the shifts of all polynomials are zero. This can be done without affecting any of the results in this work. Hence, all polynomials of depth k take values in $\mathbb{U}_{k+1}$.

Given a degree-d non-classical polynomial P, it is often useful to consider the properties of its d-th derivative. Motivated by this, we give the following definition.

Definition 2.8 (Derivative Polynomial). *Let $P : \mathbb{F}^n \to \mathbb{T}$ be a degree-d polynomial, possibly non-classical. Define the derivative polynomial $\partial P : (\mathbb{F}^n)^d \to \mathbb{T}$ by the following formula*

$$\partial P(h_1, \dots, h_d) \stackrel{\text{def}}{=} D_{h_1} \cdots D_{h_d} P(0),$$

where $h_1, \dots, h_d \in \mathbb{F}^n$.[3] *Moreover for $k < d$ define*

$$\partial_k P(x, h_1, \dots, h_k) \stackrel{\text{def}}{=} D_{h_1} \cdots D_{h_k} P(x).$$

The following lemma shows some useful properties of the derivative polynomial.

Lemma 2.9. *Let $P : \mathbb{F}^n \to \mathbb{T}$ be a degree-d (non-classical) polynomial. Then the polynomial $\partial P(h_1, \dots, h_d)$ is*

(i) multilinear: ∂P is additive in each h_i.

(ii) invariant under permutations of $h_1, \dots, h_d$.

(iii) a classical nonzero polynomial of degree d.

(iv) homogeneous: All its monomials are of degree d.

[2] Recall that $\mathbb{T}$ is an additive group. If $n \in \mathbb{Z}$ and $x \in \mathbb{T}$, then nx is shorthand for $x + \cdots + x$ if $n \geqslant 0$ and $-x - \cdots - x$ otherwise, where there are $|n|$ terms in both expressions.

[3] Notice since P is a degree d polynomial, $D_{h_1} \dots D_{h_d} P(x)$ does not depend on x and thus we have the identity $\partial P(h_1, \dots, h_d) = D_{h_1} \dots D_{h_d} P(x)$ for any choice of $x \in \mathbb{F}^n$.

Notice that by multilinear we mean additive in each direction h_i, which is not the usual use of the term "multilinear".

Proof. The proof follows by the properties of the additive derivative D_h. Multilinearity of ∂P follows from linearity of the additive derivative, namely for every function Q and directions h_1, h_2 we have the identity $D_{h_1+h_2}Q(x) = D_{h_1}Q(x) + D_{h_2}Q(x+h_1)$. The invariance under permutations of $h_1, \ldots, h_d$ is a result of commutativity of the additive derivatives. Since P is a degree-d (non-classical) polynomial, ∂P is nonzero by definition. Notice that since $D_{\underline{0}}Q \equiv 0$ for any function Q, we have $\partial P(h_1, \ldots, h_d) = 0$ if any of h_i is equal to zero. Hence every monomial of ∂P must depend on all h_i's. The properties (iii) and (iv) now follow from this and the fact that $\deg(\partial P) \leqslant d$ and thus each monomial has exactly one variable from each h_i. □

2.2 Rank of a Polynomial

We will often need to study Gowers norms of exponentials of polynomials. As we describe below if this analytic quantity is non-negligible, then there is an algebraic explanation for it: it is possible to decompose the polynomial as a function of a constant number of low-degree polynomials. To state this rigorously, let us define the notion of *rank* of a polynomial.

Definition 2.10 (Rank of a polynomial)**.** *Given a polynomial $P : \mathbb{F}^n \to \mathbb{T}$ and an integer $d > 1$, the* d-rank *of P, denoted $\mathsf{rank}_d(P)$, is defined to be the smallest integer r such that there exist polynomials $Q_1, \ldots, Q_r : \mathbb{F}^n \to \mathbb{T}$ of degree $\leqslant d-1$ and a function $\Gamma : \mathbb{T}^r \to \mathbb{T}$ satisfying $P(x) = \Gamma(Q_1(x), \ldots, Q_r(x))$. If $d = 1$, then 1-rank is defined to be ∞ if P is non-constant and 0 otherwise.*

The rank *of a polynomial $P : \mathbb{F}^n \to \mathbb{T}$ is its* $\deg(P)$*-rank. We say P is r-regular if $\mathsf{rank}(P) \geqslant r$.*

Note that for integer $\lambda \in [1, p-1]$, $\mathrm{rank}(P) = \mathrm{rank}(\lambda P)$. We also define the following weaker analytical notion of uniformity for a polynomial.

Definition 2.11 (Uniformity)**.** *Let $\varepsilon > 0$ be a real. A degree-d polynomial $P : \mathbb{F}^n \to \mathbb{T}$ is said to be ε-uniform if*

$$\|\mathsf{e}(P)\|_{U^d} < \varepsilon.$$

The following theorem of Tao and Ziegler shows that high rank polynomials have small Gowers norm.

Theorem 2.12 (Theorem 1.20 of [TZ12])**.** *For any $\varepsilon > 0$ and integer $d > 0$, there exists an integer $r(d, \varepsilon)$ such that the following is true. For any polynomial $P : \mathbb{F}^n \to \mathbb{T}$ of degree $\leqslant d$, if $\|\mathsf{e}(P)\|_{U^d} \geqslant \varepsilon$, then $\mathrm{rank}_d(P) \leqslant r$.*

This immediately implies that a regular polynomial is also uniform.

Corollary 2.13. *Let ε, d, and $r(d, \varepsilon)$ be as in Theorem 2.12. Every r-regular polynomial P of degree d is also ε-uniform.*

2.3 Polynomial Factors

A high-rank polynomial of degree d is, intuitively, a "generic" degree-d polynomial. There are no unexpected ways to decompose it into lower degree polynomials. Next, we will formalize the notion of a generic collection of polynomials. Intuitively, it should mean that there are no unexpected algebraic dependencies among the polynomials. First, we need to set up some notation.

Definition 2.14 (Factors)**.** *If X is a finite set then by a* factor $\mathcal{B}$ *we mean simply a partition of X into finitely many pieces called* atoms*.*

A function $f : X \to \mathbb{C}$ is called *$\mathcal{B}$-measurable* if it is constant on atoms of $\mathcal{B}$. For any function $f : X \to \mathbb{C}$, we may define the conditional expectation

$$\mathbf{E}[f|\mathcal{B}](x) = \mathop{\mathbf{E}}_{y \in \mathcal{B}(x)}[f(y)],$$

where $\mathcal{B}(x)$ is the unique atom in $\mathcal{B}$ that contains x. Note that $\mathbf{E}[f|\mathcal{B}]$ is $\mathcal{B}$-measurable.

A finite collection of functions $\phi_1, \ldots, \phi_C$ from X to some other space Y naturally define a factor $\mathcal{B} = \mathcal{B}_{\phi_1,\ldots,\phi_C}$ whose atoms are sets of the form $\{x : (\phi_1(x), \ldots, \phi_C(x)) = (y_1, \ldots, y_C)\}$ for some $(y_1, \ldots, y_C) \in Y^C$. By an abuse of notation we also use $\mathcal{B}$ to denote the map $x \mapsto (\phi_1(x), \ldots, \phi_C(x))$, thus also identifying the atom containing x with $(\phi_1(x), \ldots, \phi_C(x))$.

Definition 2.15 (Polynomial factors)**.** *If $P_1, \ldots, P_C : \mathbb{F}^n \to \mathbb{T}$ is a sequence of polynomials, then the factor $\mathcal{B}_{P_1,\ldots,P_C}$ is called a* polynomial factor*.*

The *complexity* of $\mathcal{B}$, denoted $|\mathcal{B}| := C$, is the number of defining polynomials. The *degree* of $\mathcal{B}$ is the maximum degree among its defining polynomials $P_1, \ldots, P_C$. If $P_1, \ldots, P_C$ are of depths $k_1, \ldots, k_C$, respectively, then the number of atoms of $\mathcal{B}$ is at most $\prod_{i=1}^{C} p^{k_i+1}$.

Definition 2.16 (Rank and Regularity)**.** *A polynomial factor $\mathcal{B}$ defined by a sequence of polynomials $P_1, \ldots, P_C : \mathbb{F}^n \to \mathbb{T}$ with respective depths $k_1, \ldots, k_C$ is said to have rank r if r is the least integer for which there exists $(\lambda_1, \ldots, \lambda_C) \in \mathbb{Z}^C$, with $(\lambda_1 \mod p^{k_1+1}, \ldots, \lambda_C \mod p^{k_C+1}) \neq 0^C$, such that $\mathrm{rank}_d(\sum_{i=1}^{C} \lambda_i P_i) \leqslant r$, where $d = \max_i \deg(\lambda_i P_i)$.*

Given a polynomial factor $\mathcal{B}$ and a function $r : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$, we say that $\mathcal{B}$ is r-regular if $\mathcal{B}$ is of rank larger than $r(|\mathcal{B}|)$.

Notice that by the above definition of rank for a degree-d polynomial P of depth k we have

$$\mathrm{rank}(\{P\}) = \min\left\{\mathrm{rank}_d(P), \mathrm{rank}_{d-(p-1)}(pP), \ldots, \mathrm{rank}_{d-k(p-1)}(p^k P)\right\}.$$

We also define the following weaker analytical notion of uniformity for a factor along the same lines as Definition 2.11.

Definition 2.17 (Uniform Factor)**.** *Let $\varepsilon > 0$ be a real. A polynomial factor $\mathcal{B}$ defined by a sequence of polynomials $P_1, \ldots, P_C : \mathbb{F}^n \to \mathbb{T}$ with respective depths $k_1, \ldots, k_C$ is said to be ε-uniform if for every collection $(\lambda_1, \ldots, \lambda_C) \in \mathbb{Z}^C$, with $(\lambda_1 \mod p^{k_1+1}, \ldots, \lambda_C \mod p^{k^C+1}) \neq 0^C$*

$$\left\| \mathsf{e}\left(\sum_i \lambda_i P_i\right) \right\|_{U^d} < \varepsilon,$$

where $d = \max_i \deg(\lambda_i P_i)$.

Remark 2.18. *Similar to Corollary 2.13 it also follows from Theorem 2.12 that an r-regular degree-d factor $\mathcal{B}$ is also ε-uniform when $r = r(d, \varepsilon)$ is as in Theorem 2.12.*

2.3.1 Regularization of Factors

Due to the generic properties of regular factors, it is often useful to *refine* a given polynomial factor to a regular one [TZ12, BFL13, BFH+13]. We will first formally define what we mean by refining a polynomial factor.

Definition 2.19 (Refinement)**.** *A factor $\mathcal{B}'$ is called a* refinement *of $\mathcal{B}$, and denoted $\mathcal{B}' \succeq \mathcal{B}$, if the induced partition by $\mathcal{B}'$ is a combinatorial refinement of the partition induced by $\mathcal{B}$. In other words, if for every $x, y \in \mathbb{F}^n$, $\mathcal{B}'(x) = \mathcal{B}'(y)$ implies $\mathcal{B}(x) = \mathcal{B}(y)$. We will write $\mathcal{B} \succeq_{syn} \mathcal{B}'$, if the polynomials defining $\mathcal{B}'$ extend that of $\mathcal{B}$.*

The following lemma from [BFH+13] which uses a regularization theorem of [TZ12] allows one to regularize a given factor to any desired regularity.

Lemma 2.20 (Polynomial Regularity Lemma [BFH+13])**.** *Let $r : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ be a non-decreasing function and $d > 0$ be an integer. Then, there is a function $C_{2.20}^{(r,d)} : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ such that the following is true. Suppose $\mathcal{B}$ is a factor defined by polynomials $P_1, \ldots, P_C : \mathbb{F}^n \to \mathbb{T}$ of degree at most d. Then, there is an r-regular factor $\mathcal{B}'$ consisting of polynomials $Q_1, \ldots, Q_{C'} : \mathbb{F}^n \to \mathbb{T}$ of degree $\leqslant d$ such that $\mathcal{B}' \succeq \mathcal{B}$ and $C' \leqslant C_{2.20}^{(r,d)}(C)$.*

2.4 Decomposition Theorems

An important application of the inverse theorems are the "decomposition theorems" [Gow10, Tao07, GT09]. These theorems allow one to express a given function f with certain properties as a sum $\sum_{i=1}^{k} g_i$, where each g_i has certain desired structural properties. We refer the interested reader to [Gow10] and [Gre07] for a detailed discussion of this subject. The following decomposition theorem is a consequence of an inverse theorem for Gowers norms ([TZ12, Theorem 1.11]).

Theorem 2.21 (Strong Decomposition Theorem for Multiple Functions)**.** *Let $m, d \geqslant 1$ be integers, $\delta > 0$ a parameter, and let $r : \mathbb{N} \to \mathbb{N}$ be an arbitrary growth function. Given any functions $f_1, \ldots, f_m : \mathbb{F}^n \to \mathbb{D}$, there exists a decomposition*

$$f_i = g_i + h_i,$$

such that for every $1 \leqslant i \leqslant m$,

1. *$g_i = \mathbf{E}[f_i|\mathcal{B}]$, where $\mathcal{B}$ is an r-regular polynomial factor of degree at most d and complexity $C \leqslant C_{\max}(p, m, d, \delta, r(\cdot))$,*
2. *$\|h_i\|_{U^{d+1}} \leqslant \delta$.*

3 Main Results

3.1 Homogeneous Polynomials

Recall that a classical polynomial is called homogeneous if all of its monomials are of the same degree. Trivially a homogeneous classical polynomial $P(x)$ satisfies $P(cx) = |c|^d P(x)$ for every $c \in \mathbb{F}$. We will use this property to define the class of non-classical homogeneous polynomials.

Definition 3.1 (Homogeneity). *A (non-classical) polynomial $P : \mathbb{F}^n \to \mathbb{T}$ is called homogeneous if for every $c \in \mathbb{F}$ there exists a $\sigma_c \in \mathbb{Z}$ such that $P(cx) = \sigma_c P(x) \mod 1$ for all x.*

Remark 3.2. *It is not difficult to see that $P(cx) = \sigma_c P(x) \mod 1$ implies that $\sigma_c = |c|^{\deg(P)} \mod p$, a property that we will use later. Indeed for $d = \deg(P)$, we have $0 = \partial_d(P(cx) - \sigma_c P(x)) = (|c|^d - \sigma_c)\partial_d P(x) \mod 1$. This, since $\partial_d P(x)$ is a nonzero degree-d classical polynomial, implies $\sigma_c = |c|^d \mod p$.*

Notice that for a polynomial P to be homogeneous it suffices that there exists $\sigma \in \mathbb{Z}$ for which $P(\zeta x) = \sigma P(x) \mod 1$, where ζ is a generator of $\mathbb{F}^*$. If P has depth k, then we can assume that $\sigma \in \mathbb{Z}_{p^{k+1}}$, as $p^{k+1}P \equiv 0$. The following lemma shows that σ is uniquely determined for all homogeneous polynomials of degree d and depth k. Henceforth, we will denote this unique value by $\sigma(d, k)$.

Lemma 3.3. *For every d and k, there is a unique $\sigma = \sigma(d, k) \in \mathbb{Z}_{p^{k+1}}$, such that for every homogeneous polynomial P of degree d and depth k, $P(\zeta x) = |\sigma| P(x) \mod 1$, where $|\cdot|$ is the natural map from $\mathbb{Z}_{p^{k+1}}$ to $\{0, 1, \ldots, p^{k+1} - 1\} \subset \mathbb{Z}$.*

Proof. Let P be a homogeneous polynomial of degree d and depth k, and let $\sigma \in \mathbb{Z}_{p^{k+1}}$ be such that $P(\zeta x) = |\sigma| P(x) \mod 1$. By Remark 3.2 we know that $|\sigma| = |\zeta|^d \mod p$. We also observe that $P(x) = P(\zeta^{p-1}x) = |\sigma|^{p-1} P(x) \mod 1$ from which it follows that $\sigma^{p-1} = 1$. We claim that $\sigma \in \mathbb{Z}_{p^{k+1}}$ is uniquely determined by the two properties

i. $|\sigma| = |c|^d \mod p$, and

ii. $\sigma^{p-1} = 1$.

Suppose to the contrary that there are two nonzero values $\sigma_1, \sigma_2 \in \mathbb{Z}_{p^{k+1}}$ that satisfy the above two properties, and choose $t \in \mathbb{Z}_{p^{k+1}}$ such that $\sigma_1 = t\sigma_2$. It follows from (i) that $t = 1 \mod p$ and from (ii) that $t^{p-1} = 1$. We will show that $t = 1$ is the only possible such value in $\mathbb{Z}_{p^{k+1}}$.

Let $a_1, \ldots, a_{p^k} \in \mathbb{Z}_{p^{k+1}}$ be all the possible solutions to $x = 1 \mod p$ in $\mathbb{Z}_{p^{k+1}}$. Note that $ta_1, \ldots, ta_{p^k}$ is just a permutation of the first sequence and thus

$$t^{p^k} \prod a_i = \prod a_i.$$

Consequently $t^{p^k} = 1$, which combined with $t^p = t$ implies $t = 1$. □

Lemma 2.6 allows us to express every (non-classical) polynomial as a linear span of monomials of the form $\frac{|x_1|^{d_1} \cdots |x_n|^{d_n}}{p^{k+1}}$. Unfortunately, unlike in the classical case, these monomials are not necessarily homogeneous, and for some applications it is important to express a polynomial as a linear span of homogeneous polynomials. We show that this is possible as homogeneous multivariate (non-classical) polynomials linearly span the space of multivariate (non-classical) polynomials. We will present the proof of this theorem in Section 4.1.

Theorem 3.4. *There is a basis for* $\mathrm{Poly}(\mathbb{F}^n \to \mathbb{T})$ *consisting only of homogeneous multivariate polynomials.*

Theorem 3.4 allows us to make the extra assumption in the strong decomposition theorem (Theorem 2.21) that the resulting polynomial factor $\mathcal{B}$ consists only of homogeneous polynomials.

Corollary 3.5. *Let $d \geqslant 1$ be an integer, $\delta > 0$ a parameter, and let $r : \mathbb{N} \to \mathbb{N}$ be an arbitrary growth function. Given any functions $f_1, \ldots, f_m : \mathbb{F}^n \to \mathbb{D}$, there exists a decomposition*

$$f_i = g_i + h_i,$$

such that or every $1 \leqslant i \leqslant m$,

1. *$g_i = \mathbf{E}[f_i|\mathcal{B}]$, where $\mathcal{B}$ is an r-regular polynomial factor of degree at most d and complexity $C \leqslant C_{\max}(p, d, \delta, r(\cdot))$, moreover $\mathcal{B}$ only consists of homogeneous polynomials.*
2. *$\|h_i\|_{U^{d+1}} \leqslant \delta$.*

3.2 Strong Near-Orthogonality

As mentioned in the introduction the main result of this paper is a new near-orthogonality result for polynomial factors of high rank. Such a statement was proved in [GT09, TZ12] for systems of linear forms corresponding to repeated derivatives or equivalently Gowers norms, in [HL11] for the case when the field is of high characteristic but with arbitrary system of linear forms and in [BFH+13] for systems of affine linear forms. In Theorem 3.7 we establish the near-orthogonality over any arbitrary system of linear forms. Before stating this theorem we need to introduce the notion of consistency.

Definition 3.6 (Consistency)**.** *Let $\mathcal{L} = \{L_1, \ldots, L_m\}$ be a system of linear forms. A vector $(\beta_1, \ldots, \beta_m) \in \mathbb{T}^m$ is said to be (d, k)-*consistent *with $\mathcal{L}$ if there exists a homogeneous polynomial P of degree d and depth k and a point X such that $P(L_i(X)) = \beta_i$ for every $i \in [m]$. Let $\Phi_{d,k}(\mathcal{L})$ denote the set of all such vectors.*

It is immediate from the definition that $\Phi_{d,k}(\mathcal{L}) \subseteq \mathbb{U}^m_{k+1}$ is a subgroup of $\mathbb{T}^m$, or more specifically, a subgroup of $\mathbb{U}^m_{k+1}$. Let

$$\Phi_{d,k}(\mathcal{L})^\perp := \left\{(\lambda_1, \ldots, \lambda_m) \in \mathbb{Z}^m \ : \ \forall (\beta_1, \ldots, \beta_m) \in \Phi_{d,k}(\mathcal{L}), \ \sum \lambda_i \beta_i = 0\right\}.$$

Equivalently $\Phi_{d,k}(\mathcal{L})^\perp$ is the set of all $(\lambda_1, \ldots, \lambda_m) \in \mathbb{Z}^m$ such that $\sum_{i=1}^m \lambda_i P(L_i(X)) \equiv 0$ for every homogeneous polynomial P of degree d and depth k.

Theorem 3.7 (Near Orthogonality over Linear Forms)**.** *Let $L_1, \ldots, L_m$ be linear forms on ℓ variables and let $\mathcal{B} = (P_1, \ldots, P_C)$ be an ε-uniform polynomial factor for some $\varepsilon \in (0, 1]$ defined only by homogeneous polynomials. For every tuple Λ of integers $(\lambda_{i,j})_{i\in[C],j\in[m]}$, define $P_\Lambda : (\mathbb{F}^n)^\ell \to \mathbb{T}$ as*

$$P_\Lambda(X) = \sum_{i\in[C],j\in[m]} \lambda_{i,j} P_i(L_j(X)).$$

Then one of the following two statements holds:

- *$P_\Lambda \equiv 0$.*
- *P_Λ is non-constant and $\left|\mathbf{E}_{X\in(\mathbb{F}^n)^\ell}[\mathsf{e}(P_\Lambda)]\right| < \varepsilon$.*

Furthermore $P_\Lambda \equiv 0$ if and only if for every $i \in [C]$, we have $(\lambda_{i,j})_{j\in[m]} \in \Phi_{d_i,k_i}(\mathcal{L})^\perp$ where d_i, k_i are the degree and depth of P_i, respectively.

We will present the proof of Theorem 3.7 in Section 4.2.

Remark 3.8. *By Corollary 2.13 the assumption of ε-uniformity in Theorem 3.7 is satisfied for every factor $\mathcal{B}$ of rank at least $r_{2.12}(d,\varepsilon)$. However, we would like to point out that in Theorem 3.7 by using the assumption of ε-uniformity instead of the assumption of high rank, we are able to achieve the quantitative bound of ε on the bias of P_Λ.*

Remark 3.9. *In Theorem 3.7 in the second case where P_Λ is non-constant, it is possible to deduce a more general statement that $\|\mathsf{e}(P_\Lambda)\|_{U^t}^{2^t} < \varepsilon$ for every $t \leqslant \deg(P_\Lambda)$. Indeed assume that P_Λ is non-constant, and consider the derivative*

$$D_{Y_1} \dots D_{Y_t} P_\Lambda(X) = \sum_{i\in[C],j\in[m]} \lambda_{i,j} \sum_{S\subseteq[t]} (-1)^{|S|} P_i(L_j(X + \sum_{r\in S} Y_r)), \tag{3}$$

where $Y_i := (y_{i,1},\dots,y_{i,\ell}) \in (\mathbb{F}^n)^\ell$. Notice that for every choice of $j \in [m]$ and $S \subseteq [t]$, $L_j(X + \sum_{r\in S} Y_r)$ is an application of a linear form on the vector $(x_1,\dots,x_\ell,y_{1,1},\dots,y_{1,\ell},\dots,y_{t,1},\dots,y_{t,\ell}) \in (\mathbb{F}^n)^{(t+1)\ell}$, and since by Lemma 2.9 the polynomial ∂P_Λ is nonzero, Theorem 3.7 implies

$$\|\mathsf{e}(P_\Lambda)\|_{U^t}^{2^t} = \mathop{\mathbf{E}}_{h_1,\dots,h_t} [\mathsf{e}(\partial_t P_\Lambda(h_1,\dots,h_t))] \leqslant \varepsilon.$$

It is well-known that statements similar to that of Theorem 3.7 imply "near-equidistributions" of the joint distribution of the polynomials applied to linear forms. Consider a highly uniform polynomial factor of degree $d > 0$, defined by a tuple of homogeneous polynomials $P_1,\dots,P_C : \mathbb{F}^n \to \mathbb{T}$ with respective degrees $d_1,\dots,d_C$ and depths $k_1,\dots,k_C$, and let $\mathcal{L} = (L_1,\dots,L_m)$ be a collection of linear forms on ℓ variables. As we mentioned earlier, we are interested in the distribution of the random matrix

$$\begin{pmatrix} P_1(L_1(X)) & P_2(L_1(X)) & \dots & P_C(L_1(X)) \\ P_1(L_2(X)) & P_2(L_2(X)) & \dots & P_C(L_2(X)) \\ \vdots & & & \vdots \\ P_1(L_m(X)) & P_2(L_m(X)) & \dots & P_C(L_m(X)) \end{pmatrix}, \tag{4}$$

where X is the uniform random variable taking values in $(\mathbb{F}^n)^\ell$. Note that by the definition of consistency, for every $1 \leqslant i \leqslant C$, the i-th column of this matrix must belong to $\Phi_{d_i,k_i}(\mathcal{L})$. Theorem 3.10 below says that Equation (4) is "almost" uniformly distributed over the set of all matrices satisfying this condition. The proof of Theorem 3.10 is standard and is identical to the proof of [BFH+13, Theorem 3.10] with the only difference that it uses Theorem 3.7 instead of the weaker near-orthogonality theorem of [BFH+13].

Theorem 3.10 (Near-equidistribution)**.** *Given $\varepsilon > 0$, let $\mathcal{B}$ be an ε-uniform polynomial factor of degree $d > 0$ and complexity C, that is defined by a tuple of homogeneous polynomials $P_1,\dots,P_C : \mathbb{F}^n \to \mathbb{T}$ having respective degrees $d_1,\dots,d_C$ and depths $k_1,\dots,k_C$. Let $\mathcal{L} = (L_1,\dots,L_m)$ be a collection of linear forms on ℓ variables.*

Suppose $(\beta_{i,j})_{i\in[C],j\in[m]} \in \mathbb{T}^{C\times m}$ is such that $(\beta_{i,1},\dots,\beta_{i,m}) \in \Phi_{d_i,k_i}(\mathcal{L})$ for every $i \in [C]$. Then

$$\mathop{\mathbf{Pr}}_{X\in(\mathbb{F}^n)^\ell} [P_i(L_j(X)) = \beta_{i,j} \ \forall i \in [C], j \in [m]] = \frac{1}{K} \pm \varepsilon,$$

where $K = \prod_{i=1}^C |\Phi_{d_i,k_i}(\mathcal{L})|$.

Proof. We have

$$\mathbf{Pr}[P_i(L_j(X)) = \beta_{i,j}\ \forall i \in [C], \forall j \in [m]] = \mathbf{E}\left[\prod_{i,j} \frac{1}{p^{k_i+1}} \sum_{\lambda_{i,j}=0}^{p^{k_i+1}-1} \mathsf{e}\left(\lambda_{i,j}\big(P_i(L_j(X)) - \beta_{i,j}\big)\right)\right]$$
$$= \left(\prod_{i\in[C]} p^{-(k_i+1)}\right)^m \sum_{(\lambda_{i,j})} \mathsf{e}\left(-\sum_{i,j}\lambda_{i,j}\beta_{i,j}\right) \mathbf{E}\left[\mathsf{e}\left(\sum_{i\in[C],j\in[m]} \lambda_{i,j}P_i(L_j(X))\right)\right],$$

where the outer sum is over $(\lambda_{i,j})_{i\in[C],j\in[m]}$ with $\lambda_{i,j} \in [0, p^{k_i+1}-1]$. Let $\Lambda_i = \Phi_{d_i,k_i}(\mathcal{L})^\perp \cap [0, p^{k+1}-1]^m$, and note that $|\Lambda_i||\Phi_{d_i,k_i}| = p^{m(k_i+1)}$. Since $(\beta_{i,1},\ldots,\beta_{i,m}) \in \Phi_{d_i,k_i}(\mathcal{L})$ for every $i \in [C]$, it follows that $\sum_{i,j}\lambda_{i,j}\beta_{i,j} = 0$ if $(\lambda_{i,1},\ldots,\lambda_{i,m}) \in \Lambda_i$ for all $i \in [C]$. If the latter holds, then the expected value in the above expression is 0, and otherwise by Theorem 3.7, it is bounded by ε. Hence the above expression can be approximated by

$$p^{-m\sum_{i=1}^C (k_i+1)} \cdot \left(\prod_{i=1}^C |\Lambda_i| \ \pm\ \varepsilon p^{m\sum_{i=1}^C (k_i+1)}\right) = \frac{1}{K} \pm \varepsilon.$$

□

3.3 On a Theorem of Gowers and Wolf

Let A be a subset of $\mathbb{F}^n$ with the indicator function $\mathbf{1}_A : \mathbb{F}^n \to \{0,1\}$. As mentioned in the introduction, Equation (1) equals the probability that $L_1(X),\ldots,L_m(X)$ all fall in A, where $X \in (\mathbb{F}^n)^k$ is chosen uniformly at random. Roughly speaking, we say $A \subseteq \mathbb{F}^n$ is *pseudorandom* with regards to $\mathcal{L}$ if

$$\mathop{\mathbf{E}}_X\left[\prod_{i=1}^m \mathbf{1}_A(L_i(X))\right] \approx \left(\frac{|A|}{p^n}\right)^m;$$

That is if the probability that all $L_1(X),\ldots,L_m(X)$ fall in A is close to what we would expect if A was a random subset of $\mathbb{F}^n$ of cardinality $|A|$. Let $\alpha := |A|/p^n$ be the density of A, and define $f := \mathbf{1}_A - \alpha$. We have

$$\mathop{\mathbf{E}}_X\left[\prod_{i=1}^m \mathbf{1}_A(L_i(X))\right] = \mathop{\mathbf{E}}_X\left[\prod_{i=1}^m (\alpha + f(L_i(X)))\right] = \alpha^m + \sum_{S\subseteq[m],S\neq\emptyset} \alpha^{m-|S|} \mathop{\mathbf{E}}_X\left[\prod_{i\in S} f(L_i(X))\right].$$

Therefore, a sufficient condition for A to be pseudorandom with regards to $\mathcal{L}$ is that $\mathbf{E}_X\left[\prod_{i\in S} f(L_i(X))\right]$ is negligible for all nonempty subsets $S \subseteq [m]$. Green and Tao [GT10b] showed that a sufficient condition for this to occur is that $\|f\|_{U^{s+1}}$ is small enough, where s is the *Cauchy-Schwarz complexity* of the system of linear forms.

Definition 3.11 (Cauchy-Schwarz complexity [GT10b])**.** *Let $\mathcal{L} = \{L_1,\ldots,L_m\}$ be a system of linear forms. The* Cauchy-Schwarz complexity *of $\mathcal{L}$ is the minimal s such that the following holds. For every $1 \leqslant i \leqslant m$, we can partition $\{L_j\}_{j\in[m]\setminus\{i\}}$ into $s+1$ subsets, such that L_i does not belong to the linear span of any of the subsets.*

The reason for the term *Cauchy-Schwarz complexity* is the following lemma due to Green and Tao [GT10b] whose proof is based on a clever iterative application of the Cauchy-Schwarz inequality.

Lemma 3.12 ([GT10b], See also [GW10, Theorem 2.3]). *Let $f_1, \dots, f_m : \mathbb{F} \to \mathbb{D}$. Let $\mathcal{L} = \{L_1, \dots, L_m\}$ be a system of m linear forms in ℓ variables of Cauchy-Schwarz complexity s. Then*

$$\left| \mathop{\mathbf{E}}_{X \in (\mathbb{F}^n)^\ell} \left[\prod_{i=1}^{m} f_i(L_i(X)) \right] \right| \leqslant \min_{1 \leqslant i \leqslant m} \|f_i\|_{U^{s+1}}.$$

Note that the Cauchy-Schwarz complexity of any system of m linear forms in which any two linear forms are linearly independent (i.e. one is not a multiple of the other) is at most $m-2$, since we can always partition $\{L_j\}_{j \in [m] \setminus \{i\}}$ into the $m-1$ singleton subsets.

The Cauchy-Schwarz complexity of $\mathcal{L}$ gives an upper bound on s, such that if $\|f\|_{U^{s+1}}$ is small enough for some function $f : \mathbb{F}^n \to \mathbb{D}$, then f is pseudorandom with regards to $\mathcal{L}$. Gowers and Wolf [GW10] defined the *true complexity* of a system of linear forms as the minimal s such that the above condition holds for all $f : \mathbb{F}^n \to \mathbb{D}$.

Definition 3.13 (True complexity [GW10]). *Let $\mathcal{L} = \{L_1, \dots, L_m\}$ be a system of linear forms over $\mathbb{F}$. The true complexity of $\mathcal{L}$ is the smallest $d \in \mathbb{N}$ with the following property. For every $\varepsilon > 0$, there exists $\delta > 0$ such that if $f : \mathbb{F}^n \to \mathbb{D}$ is any function with $\|f\|_{U^{d+1}} \leqslant \delta$, then*

$$\left| \mathop{\mathbf{E}}_{X \in (\mathbb{F}^n)^k} \left[\prod_{i=1}^{m} f(L_i(X)) \right] \right| \leqslant \varepsilon.$$

An obvious bound on the true complexity is the Cauchy-Schwarz complexity of the system. However, there are cases where this is not tight. Gowers and Wolf conjectured that the true complexity of a system of linear forms can be characterized by a simple linear algebraic condition. Namely, that it is equal to the smallest $d \geqslant 1$ such that $L_1^{d+1}, \dots, L_m^{d+1}$ are linearly independent where the d-th tensor power of a linear form L is defined as

$$L^d = \left(\prod_{j=1}^{d} \lambda_{i_j} : i_1, \dots, i_d \in [k] \right) \in \mathbb{F}^{k^d}.$$

Later in [GW11, Theorem 6.1] they verified their conjecture in the case where $|\mathbb{F}|$ is sufficiently large; more precisely when $|\mathbb{F}|$ is at least the Cauchy-Schwarz complexity of the system of linear form. In this paper we verify the Gowers-Wolf conjecture in full generality by proving the following stronger theorem.

Theorem 3.14. *Let $\mathcal{L} = \{L_1, \dots, L_m\}$ be a system of linear forms. Assume that L_1^{d+1} is not in the linear span of $L_2^{d+1}, \dots, L_m^{d+1}$. For every $\varepsilon > 0$, there exists $\delta > 0$ such that for any collection of functions $f_1, \dots, f_m : \mathbb{F}^n \to \mathbb{D}$ with $\|f_1\|_{U^{d+1}} \leqslant \delta$, we have*

$$\left| \mathop{\mathbf{E}}_{X \in (\mathbb{F}^n)^k} \left[\prod_{i=1}^{m} f_i(L_i(X)) \right] \right| \leqslant \varepsilon.$$

Theorem 3.14 was conjectured in [GW11], and left open even in the case of large $|\mathbb{F}|$. In [HL11], a partial near-orthogonality result is proved and used to prove Theorem 3.14 in the case where $|\mathbb{F}|$ is greater or equal to the Cauchy-Schwarz complexity of the system of linear form. In this paper, our full near-orthogonality result allows us to establish this theorem in its full generality. We will present the proof of Theorem 3.14 in Section 4.3. The following corollary to Theorem 3.14 is very useful when combined with the decomposition theorems such as Theorem 2.21.

Corollary 3.15. *Let $\mathcal{L} = \{L_1, \ldots, L_m\}$ be a system of linear forms. Assume that $L_1^{d+1}, \ldots, L_m^{d+1}$ are linearly independent. For every $\varepsilon > 0$, there exists $\delta > 0$ such that for any functions $f_1, \ldots, f_m, g_1, \ldots, g_m : \mathbb{F}^n \to \mathbb{D}$ with $\|f_i - g_i\|_{U^{d+1}} \leqslant \delta$, we have*

$$\left| \mathop{\mathbf{E}}_X \left[\prod_{i=1}^m f_i(L_i(X)) \right] - \mathop{\mathbf{E}}_X \left[\prod_{i=1}^m g_i(L_i(X)) \right] \right| \leqslant \varepsilon,$$

Proof. Choosing $\delta = \delta(\varepsilon')$ as in Theorem 3.14 for $\varepsilon' := \varepsilon/m$, we have

$$\begin{aligned}
\left| \mathop{\mathbf{E}}_X \left[\prod_{i=1}^m f_i(L_i(X)) \right] - \mathop{\mathbf{E}}_X \left[\prod_{i=1}^m g_i(L_i(X)) \right] \right| &= \left| \sum_{i=1}^m \mathop{\mathbf{E}}_X \left[(f_i - g_i)(L_i(X)) \cdot \prod_{j=1}^{i-1} g_j(L_j(X)) \cdot \prod_{j=i+1}^m f_j(L_j(X)) \right] \right| \\
&\leqslant \sum_{i=1}^m \left| \mathop{\mathbf{E}}_X \left[(f_i - g_i)(L_i(X)) \cdot \prod_{j=1}^{i-1} g_j(L_j(X)) \cdot \prod_{j=i+1}^m f_j(L_j(X)) \right] \right| \\
&\leqslant m \cdot \delta \leqslant \varepsilon,
\end{aligned}$$

where the second inequality follows from Theorem 3.14 since $\|f_i - g_i\|_{U^{d+1}} \leqslant \delta$ and L_i^{d+1} is not in the linear span of $\{L_j^{d+1}\}_{j \in [m] \setminus \{i\}}$. □

4 Main Proofs

In this section we will present the proofs of Theorem 3.4, Theorem 3.7, and Theorem 3.14.

4.1 Homogeneity: Proof of Theorem 3.4

Theorem 3.4 (restated). *There is a basis for* $\mathrm{Poly}(\mathbb{F}^n \to \mathbb{T})$ *consisting only of homogeneous multivariate polynomials.*

To simplify the notation, in this section we will omit writing "mod 1" in the description of the defined non-classical polynomials. We start by proving the following simple observation.

Claim 4.1. *Let $P : \mathbb{F} \to \mathbb{T}$ be a univariate polynomial of degree d. Then for every $c \in \mathbb{F} \setminus \{0\}$,*

$$\deg \Big(P(cx) - |c|^d P(x) \Big) < d.$$

Proof. By Lemma 2.6 it suffices to prove the claim for a monomial $q(x) := \frac{|x|^s}{p^{k+1}}$ with $k(p-1)+s = d$. Note that $q(cx) - |c|^d q(x)$ takes values in $\frac{1}{p^k}\mathbb{Z}/\mathbb{Z}$ as $|c|^s - |c|^d$ is divisible by p. Hence

$$\deg \Big(q(cx) - |c|^d q(x) \Big) \leqslant (p-1)(k-1) < d. \tag{5}$$

□

Remark 4.2. *It is not difficult to show that the above claim holds for any multivariate polynomial $P : \mathbb{F}^n \to \mathbb{T}$. However, since the univariate case suffices for our purpose, we do not prove the general case.*

First we prove Theorem 3.4 for univariate polynomials.

Lemma 4.3. *There is a basis of homogeneous univariate polynomials for* $\mathrm{Poly}(\mathbb{F} \to \mathbb{T})$.

Proof. We will prove by induction on d that there is a basis $\{h_1, \ldots, h_d\}$ of homogeneous univariate polynomials for $\mathrm{Poly}_{\leqslant d}(\mathbb{F} \to \mathbb{T})$ for every d. Let ζ be a fixed generator of $\mathbb{F}^*$. For any degree $d > 0$, we will build a degree-d homogeneous polynomial $h_d(x)$ such that $h_d(\zeta x) = \sigma_d h_d(x)$ for some integer σ_d. The base case of $d \leqslant p-1$ is trivial as $\mathrm{Poly}_{\leqslant p-1}(\mathbb{F} \to \mathbb{T})$ consists of only classical polynomials, and those are spanned by $h_0(x) := \frac{1}{p}, h_1(x) := \frac{|x|}{p}, \ldots, h_{p-1}(x) := \frac{|x|^{p-1}}{p}$. Now suppose that $d = s + (p-1)(k-1)$ with $0 < s \leqslant p-1$, and $k > 1$. It suffices to show that the degree-d monomial $\frac{|x|^s}{p^k}$ can be expressed as a linear combination of homogeneous polynomials. Consider the function

$$f(x) := \frac{|\zeta x|^s}{p^k} - \frac{|\zeta|^s |x|^s}{p^k}.$$

Claim 4.1 implies that $\deg(f) < d$. Using the induction hypothesis, we can express $f(x)$ as a linear combination of $\frac{|x|^s}{p^\ell}$ for $\ell = 0, \ldots, k-1$, and h_e for $e < d$ with $e \neq s \mod (p-1)$:

$$f(x) = \sum_{\ell=1}^{k-1} a_\ell \frac{|x|^s}{p^\ell} + \sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} b_e h_e(x).$$

Set $A := |\zeta|^s + \sum_{\ell=1}^{k-1} a_\ell p^{k-\ell}$, so that

$$\frac{|\zeta x|^s}{p^k} - A\frac{|x|^s}{p^k} = \sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} b_e h_e(x). \tag{6}$$

By the induction hypothesis, for $e < d$, $h_e(\zeta x) = \sigma_e h(x)$ where $\sigma_e = |\zeta|^e \mod p$, and thus as $A = |\zeta|^s \mod p$, we have $\sigma_e \neq A \mod p$ when $e \neq s \mod (p-1)$. Consequently,

$$\begin{aligned}\sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} b_e h_e(x) &= \sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} \frac{b_e}{\sigma_e - A}(\sigma_e - A) h_e(x) \\ &= \sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} \frac{b_e}{\sigma_e - A}(h_e(\zeta x) - A h_e(x)).\end{aligned}$$

Combing this with (6) we conclude that

$$h_d(x) := \frac{|x|^s}{p^k} - \sum_{\substack{e<d, \\ e \neq d \mod (p-1)}} \frac{b_e}{\sigma_e - A} h_e(x),$$

satisfies

$$h_d(\zeta x) = A h_d(x).$$

□

Proof of Theorem 3.4: We will show by induction on the degree d, that every degree d monomial can be written as a linear combination of homogeneous polynomials. The base case of $d < p$ is trivial as such monomials are classical and thus homogeneous themselves. Consider a (non-classical) monomial $M(x_1, \ldots, x_n) = \frac{|x_1|^{s_1} \cdots |x_n|^{s_n}}{p^k}$ of degree $d = s_1 + \cdots + s_n + (p-1)(k-1)$. For every $i \in [n]$ let $g_i(x_i) := h_{s_i+(p-1)(k-1)}(x_i)$ where $h_{s_i+(p-1)(k-1)}(\cdot)$ is the homogeneous univariate polynomial from Lemma 4.3. Every g_i takes values in $\frac{1}{p^k}\mathbb{Z}/\mathbb{Z}$, and thus corresponds to a polynomial $G_i : \mathbb{F} \to \mathbb{Z}_{p^k}$. Define $F : \mathbb{F}^n \to \mathbb{Z}_{p^k}$ as

$$F(x_1, \ldots, x_n) := G_1(x_1) \cdots G_n(x_n),$$

and $f : \mathbb{F}^n \to \mathbb{T}$ as

$$f(x_1, \ldots, x_n) := \frac{F(x_1, \ldots, x_n)}{p^k}.$$

It is simple to verify that $\deg(f) = s_1 + \ldots + s_n + (p-1)(k-1) = d$, it has only one monomial of $\deg(f)$, which is $\frac{|x_1|^{s_1} \cdots |x_n|^{s_n}}{p^k} = M(x_1, \ldots, x_n)$, and it is homogeneous. Thus $M(x_1, \ldots, x_n) - f(x_1, \ldots, x_n)$ is of degree less than d and by the induction hypothesis can be written as a linear combination of homogeneous polynomials. □

4.2 Near-orthogonality: Proof of Theorem 3.7

Theorem 3.7 (restated). *Let $L_1, \ldots, L_m$ be linear forms on ℓ variables and let $\mathcal{B} = (P_1, \ldots, P_C)$ be an ε-uniform polynomial factor for some $\varepsilon \in (0, 1]$ defined only by homogeneous polynomials. For every tuple Λ of integers $(\lambda_{i,j})_{i\in[C],j\in[m]}$, define $P_\Lambda : (\mathbb{F}^n)^\ell \to \mathbb{T}$ as*

$$P_\Lambda(X) = \sum_{i\in[C],j\in[m]} \lambda_{i,j} P_i(L_j(X)).$$

Then one of the following two statements holds:

- *$P_\Lambda \equiv 0$.*
- *P_Λ is non-constant and $\left|\mathbf{E}_{X\in(\mathbb{F}^n)^\ell}[\mathsf{e}(P_\Lambda)]\right| < \varepsilon$.*

Furthermore $P_\Lambda \equiv 0$ if and only if for every $i \in [C]$, we have $(\lambda_{i,j})_{j\in[m]} \in \Phi_{d_i,k_i}(\mathcal{L})^\perp$ where d_i, k_i are the degree and depth of P_i, respectively.

We prove Theorem 3.7 in this section. Our proof uses similar derivative techniques as used in [BFH+13], but in order to handle the general setting we will need a few technical claims which we present first. Recall that $|L| = \sum_{i=1}^{\ell} |\lambda_i|$ for a linear form $L = (\lambda_1, \ldots, \lambda_\ell)$.

Claim 4.4. *Let $d > 0$ be an integer, and $L = (\lambda_1, \ldots, \lambda_\ell) \in \mathbb{F}^\ell$ be a linear form on ℓ variables. There exists linear forms $L_i = (\lambda_{i,1}, \ldots, \lambda_{i,\ell}) \in \mathbb{F}^\ell$ for $i = 1, \ldots, m$, and coefficients $a_1, \ldots, a_m \in \mathbb{Z}$ with $m \leqslant |\mathbb{F}|^\ell$ such that*

- *$P(L(X)) = \sum_{i=1}^{m} a_i P(L_i(X))$ for every degree-d polynomial $P : \mathbb{F}^n \to \mathbb{T}$;*
- *$|L_i| \leqslant d$ for every $i \in [m]$;*
- *$\lambda_{i,j} \leqslant \lambda_j$ for every $i \in [m]$ and $j \in [\ell]$.*

Proof. The proof proceeds by simplifying $P(L(X))$ using identities that are valid for every polynomial $P : \mathbb{F}^n \to \mathbb{T}$ of degree d.

In the case $|L| \leqslant d$ there is nothing to prove. Assume otherwise that $|L| > d$. We will use the fact that for every choice of $y_1, \ldots, y_{|L|} \in \mathbb{F}^n$,

$$\sum_{S \subseteq [|L|]} (-1)^{|S|} P\left(\sum_{i \in S} y_i\right) \equiv 0. \tag{7}$$

Let $X = (x_1, \ldots, x_\ell) \in (\mathbb{F}^n)^\ell$. Setting $|\lambda_i|$ of the vectors $y_1, \ldots, y_{|L|}$ to x_i for every $i \in [\ell]$, Equation (7) implies

$$P(L(X)) = \sum_i \alpha_i P(M_i(X)),$$

where $M_i = (\tau_{i,1}, \ldots, \tau_{i,\ell})$, $|M_i| \leqslant |L| - 1$ and for every $j \in [\ell]$, $|\tau_{i,j}| \leqslant |\lambda_j|$. Repeatedly applying the same process to every M_i with $|M_i| > d$ we arrive at the desired expansion. □

The next claim shows that we can further simplify the expression given in Claim 4.4. Let $\mathcal{L}_d \subseteq \mathbb{F}^\ell$ denote the set of nonzero linear forms L with $|L| \leqslant d$ and with the first (left-most) nonzero coefficient equal to 1, e.g. $(0, 1, 0, 2) \in \mathcal{L}_3$ but $(2, 1, 0, 0) \notin \mathcal{L}_3$.

Claim 4.5. *For any linear form $L \in \mathbb{F}^\ell$ and integer $d > 0$, there is a collection of coefficients $\{a_{M,c} \in \mathbb{Z}\}_{M \in \mathcal{L}_d, c \in \mathbb{F}^*}$ such that for every degree-d polynomial $P : \mathbb{F}^n \to \mathbb{T}$,*

$$P(L(X)) = \sum_{M \in \mathcal{L}_d, c \in \mathbb{F}^*} a_{M,c} P(cM(X)). \tag{8}$$

Proof. Similar to the proof of Claim 4.4 we simplify $P(L(X))$ using identities that are valid for every polynomial $P : \mathbb{F}^n \to \mathbb{T}$ of degree d.

We use induction on the number of nonzero entries of L. The case when L has only one nonzero entry is trivial. For the induction step, choose $c \in \mathbb{F}$ so that the leading nonzero coefficient of $L' = c \cdot L$ is equal to 1. Assume that $L' = (\lambda_1, \ldots, \lambda_\ell)$. If $|L'| \leqslant d$ we are done. Assume otherwise that $|L'| > d$. Applying Claim 4.4 for the degree-d polynomial $R(x) := P(c^{-1}x)$ and the linear form L' we can write

$$P(L(X)) = P(c^{-1}L'(X)) = \sum_i \beta_i P(c^{-1}M_i(X)), \tag{9}$$

where for every i, $M_i = (\lambda_{i,1}, \ldots, \lambda_{i,\ell})$ satisfies $|M_i| \leqslant d$, and for every $j \in [\ell]$, $\lambda_{i,j} \leqslant \lambda_j$. Let $\mathcal{I}$ denote the set of indices i such that the leading nonzero entry of M_i is one. Then

$$P(L(x)) = \sum_{i \in \mathcal{I}} \alpha_i P(c^{-1}M_i(X)) + \sum_{j \notin \mathcal{I}} \alpha_j P(c^{-1}M_j(X)). \tag{10}$$

Notice that since the leading coefficient of L' is 1, for every $j \notin \mathcal{I}$, M_j has smaller support than L and thus applying the induction hypothesis to the linear forms $c^{-1}M_j$ with $j \notin \mathcal{I}$ concludes the claim. □

Claim 4.5 applies to all polynomials of degree d. If we also specify the depth then we can obtain a stronger statement.

Claim 4.6. *For every d, k, every system of linear forms $\{L_1, \ldots, L_m\}$, and constants $\{\lambda_i \in \mathbb{Z}\}_{i\in[m]}$, there exists $\{a_M \in \mathbb{Z}\}_{M\in\mathcal{L}_d}$ such that the following is true for every (d,k)-homogeneous polynomial $P : \mathbb{F}^n \to \mathbb{T}$:*

- $\sum_{i=1}^m \lambda_i P(L_i(X)) \equiv \sum_{M\in\mathcal{L}_d} a_M P(M(X))$;
- *For every M with $a_M \neq 0$, we have $|M| \leqslant \deg(a_M P)$.*

Proof. The proof is similar to that of Claim 4.5, except that now we repeatedly apply Claim 4.5 to every term of the form $\lambda P(L(X))$ to express it as a linear combination of $P(cM(X))$ for $M \in \mathcal{L}_{\deg(\lambda P)}$ and $c \in \mathbb{F}^*$. Then we use homogeneity to replace $P(cM(X))$ with $\sigma_c P(M(X))$, where if $c = \zeta^i$ for the fixed generator $\zeta \in \mathbb{F}^*$ then $\sigma_c = \sigma(d,k)^i$. By repeating this procedure we arrive at the desired expansion. □

We are now ready for the proof of our main theorem. For a linear form $L = (\lambda_1, \ldots, \lambda_\ell)$, let $\mathrm{lc}(L)$ denote the index of its first nonzero entry, namely $\mathrm{lc}(L) \stackrel{\text{def}}{=} \min_{i:\lambda_i\neq 0} i$.

Proof of Theorem 3.7: Let d' be the degree of the factor. For every $i \in [C]$, by Claim 4.6 we have

$$\sum_{j=1}^m \lambda_{i,j} P_i(L_j(X)) = \sum_{M\in\mathcal{L}_{d'}} \lambda'_{i,M} P_i(M(X)) \tag{11}$$

for some integers $\lambda'_{i,M}$ such that $|M| \leqslant \deg(\lambda'_{i,M} P_i)$ if $\lambda'_{i,M} \neq 0$. The simplifications of Claim 4.6 depend only on the degrees and depths of the polynomials. Hence if $\lambda'_{i,M} = 0$ for all $M \in \mathcal{L}_{d'}$, then $(\lambda_{i,1}, \ldots, \lambda_{i,m}) \in \Phi(d_i, k_i)^\perp$. So to prove the theorem, it suffices to show that P_Λ has small bias if $\lambda'_{i,M} P_i \not\equiv 0$ for some $M \in \mathcal{L}_{d'}$. Suppose this is true, and thus there exists a nonempty set $\mathcal{M} \subseteq \mathcal{L}_{d'}$ such that

$$P_\Lambda(X) = \sum_{i\in[C], M\in\mathcal{M}} \lambda'_{i,M} P_i(M(X)),$$

and for every $M \in \mathcal{M}$, there is at least one index $i \in [C]$ for which $\lambda'_{i,M} P_i \neq 0$. Choose $i^* \in [C]$ and $M^* \in \mathcal{M}$ in the following manner.

- First, let $M^* \in \mathcal{M}$ be such that $\mathrm{lc}(M^*) = \min_{M\in\mathcal{M}} \mathrm{lc}(M)$, and among these, $|M^*|$ is maximal.
- Then, let $i^* \in [C]$ be such that $\deg(\lambda'_{i^*,M^*} P_{i^*})$ is maximized.

Without loss of generality assume that $i^* = 1$, $\mathrm{lc}(M^*) = 1$, and let $d := \deg(\lambda'_{1,M^*} P_1)$. We claim that if $\sum_{j\in[m]} \lambda_{1,j} P_1(L_j(X))$ is not the zero polynomial, then $\deg(P_\Lambda) \geqslant d$, and moreover P_Λ has small bias. We prove this by deriving P_Λ in specific directions in a manner that all the terms but $\lambda'_{1,M^*} P_1(M^*(X))$ vanish.

Given a vector $\alpha \in \mathbb{F}^\ell$, an element $y \in \mathbb{F}^n$, and a function $P : (\mathbb{F}^n)^\ell \to \mathbb{T}$, define the derivative of P according to the pair (α, y) as

$$D_{\alpha,y} P(x_1, \ldots, x_\ell) \stackrel{\text{def}}{=} P(x_1 + \alpha_1 y, \cdots, x_\ell + \alpha_\ell y) - P(x_1, \ldots, x_\ell). \tag{12}$$

Note that for every $M \in \mathcal{M}$,

$$\begin{aligned} D_{\alpha,y}(P_i \circ M)(x_1, \cdots, x_\ell) &= P_i(M(x_1, \ldots, x_\ell) + M(\alpha)y) - P_i(M(x_1, \ldots, x_\ell)) \\ &= (D_{\langle M,\alpha\rangle \cdot y} P_i)(M(x_1, \ldots, x_\ell)). \end{aligned}$$

Thus if α is chosen such that $\langle M, \alpha \rangle = 0$ then $D_{\alpha,y}(P_i \circ M) \equiv 0$.

Assume that $M^* = (w_1, \ldots, w_\ell)$, where $w_1 = 1$. Let $t := |M^*|$, $\alpha_1 := e_1 = (1, 0, 0, \ldots, 0) \in \mathbb{F}^\ell$, and let $\alpha_2, \ldots, \alpha_t$ be the set of all vectors of the form $(-w, 0, \ldots, 0, 1, 0, \ldots, 0)$ where 1 is in the i-th coordinate for $i \in [2, \ell]$ and $0 \leqslant w \leqslant w_i - 1$. In addition, pick $\alpha_{t+1} = \cdots = \alpha_d = e_1$.

Claim 4.7.

$$D_{\alpha_1,y_1} \cdots D_{\alpha_d,y_d} P_\Lambda(X) = \left(D_{\langle M^*,\alpha_1\rangle y_1} \cdots D_{\langle M^*,\alpha_d\rangle y_d} \sum_{\substack{i\in[C]:\\ \deg(\lambda'_{i,M^*}P_i)=d}} \lambda'_{i,1} P_i \right)(M^*(X)). \tag{13}$$

Proof. Deriving according to $(\alpha_1, y_1), \ldots, (\alpha_t, y_t)$ gives

$$D_{\alpha_1,y_1} \cdots D_{\alpha_t,y_t} P_\Lambda(X) = \left(D_{\langle M^*,\alpha_1\rangle y_1} \cdots D_{\langle M^*,\alpha_t\rangle y_t} \left(\sum_{i=1}^{C} \lambda'_{i,M^*} P_i \right) \right)(M^*(X)). \tag{14}$$

This is because for every $M = (w'_1, \ldots, w'_\ell) \in \mathcal{M} \setminus \{M^*\}$, either $w'_1 = 0$ in which case $\langle M, \alpha_1 \rangle = 0$, or otherwise $w_1 = 1$ and $|M| \leqslant t$, thus there must be an index $\xi \in [\ell]$ such that $w'_\xi < w_\xi$; By our choice of $\alpha_2, \ldots, \alpha_t$ there is e, $2 \leqslant e \leqslant t$, such that $\alpha_e = (-w'_\xi, 0, \ldots, 0, 1, 0, \ldots, 0)$ where 1 is in the ξ-th coordinate and thus $\langle M_j, \alpha_e \rangle = 0$. Now, the claim follows after additionally deriving according to $(\alpha_{t+1}, y_{t+1}), \ldots, (\alpha_d, y_d)$. □

Claim 4.7 implies that

$$\mathop{\mathbf{E}}_{\substack{y_1,\ldots,y_d,\\ x_1,\ldots,x_\ell}} \left[\mathsf{e}\left((D_{\alpha_1,y_1} \cdots D_{\alpha_d,y_d} P_\Lambda)(x_1, \ldots, x_\ell) \right) \right] = \left\| \mathsf{e}\left(\sum_{\substack{i\in[C]:\\ \deg(\lambda'_{i,1}P_i)=d}} \lambda'_{i,M^*} P_i \right) \right\|_{U^d}^{2^d} \leqslant \varepsilon^{2^d},$$

where the last inequality holds by the ε-uniformity of the polynomial factor. Now the theorem follows from the next claim from [BFH+13] which is a repeated application of the Cauchy-Schwarz inequality. We include a proof for self-containment.

Claim 4.8 ([BFH+13, Claim 3.4])**.** *For any $\alpha_1, \ldots, \alpha_d \in \mathbb{F}^\ell \setminus \{\underline{0}\}$,*

$$\mathop{\mathbf{E}}_{\substack{y_1,\ldots,y_d,\\ x_1,\ldots,x_\ell}} \left[\mathsf{e}\left((D_{\alpha_1,y_1} \cdots D_{\alpha_d,y_d} P_\Lambda)(x_1, \ldots, x_\ell) \right) \right] \geqslant \left(\left| \mathop{\mathbf{E}}_{x_1,\ldots,x_\ell} \mathsf{e}\left(P_\Lambda(x_1, \ldots, x_\ell) \right) \right| \right)^{2^d}.$$

Proof. It suffices to show that for any function $P(x_1, \ldots, x_\ell)$ and nonzero $\boldsymbol{\alpha} \in \mathbb{F}^\ell$,

$$\left| \mathop{\mathbf{E}}_{y,x_1,\ldots,x_\ell} \left[\mathsf{e}\left((D_{\boldsymbol{\alpha},y} P)(x_1, \ldots, x_\ell) \right) \right] \right| \geqslant \left| \mathop{\mathbf{E}}_{x_1,\ldots,x_\ell} \left[\mathsf{e}\left(P(x_1, \ldots, x_\ell) \right) \right] \right|^2.$$

Recall that $(D_{\boldsymbol{\alpha},y} P)(x_1, \ldots, x_\ell) = P(x_1 + \alpha_1 y, \ldots, x_\ell + \alpha_\ell y) - P(x_1, \ldots, x_\ell)$. Without loss of generality, suppose $\alpha_1 \neq 0$. We make a change of coordinates so that $\boldsymbol{\alpha}$ can be assumed to be $(1, 0, \ldots, 0)$. More precisely, define $P' : (\mathbb{F}^n)^\ell \to \mathbb{T}$ as

$$P'(x_1, \ldots, x_\ell) = P\left(x_1, \frac{x_2 + \alpha_2 x_1}{\alpha_1}, \frac{x_3 + \alpha_3 x_1}{\alpha_1}, \ldots, \frac{x_\ell + \alpha_\ell x_1}{\alpha_1} \right),$$

so that $P(x_1,\dots,x_\ell) = P'(x_1,\alpha_1x_2-\alpha_2x_1,\alpha_1x_3-\alpha_3x_1,\dots,\alpha_1x_\ell-\alpha_\ell x_1)$, and thus $(D_{\boldsymbol{\alpha},y}P)(x_1,\dots,x_\ell) = P'(x_1+\alpha_1y,\alpha_1x_2-\alpha_2x_1,\dots,\alpha_1x_\ell-\alpha_\ell x_1) - P'(x_1,\alpha_1x_2-\alpha_2x_1,\dots,\alpha_1x_\ell-\alpha_\ell x_1)$. Therefore

$$\begin{aligned}
&\left|\mathop{\mathbf{E}}_{y,x_1,\dots,x_\ell}[\mathsf{e}\left((D_{\boldsymbol{\alpha},y}P)(x_1,\dots,x_\ell)\right)]\right| \\
&= \left|\mathop{\mathbf{E}}_{y,x_1,\dots,x_\ell}[\mathsf{e}\left(P'(x_1+\alpha_1y,\alpha_1x_2-\alpha_2x_1,\dots,\alpha_1x_\ell-\alpha_\ell x_1) - P'(x_1,\alpha_1x_2-\alpha_2x_1,\dots,\alpha_1x_\ell-\alpha_\ell x_1)\right)]\right| \\
&= \left|\mathop{\mathbf{E}}_{y,x_1,\dots,x_\ell}[\mathsf{e}\left(P'(x_1+\alpha_1y,x_2,\dots,x_\ell) - P'(x_1,x_2,\dots,x_\ell)\right)]\right| = \mathop{\mathbf{E}}_{x_2,\dots,x_\ell}\left|\mathop{\mathbf{E}}_{x_1}[\mathsf{e}\left(P'(x_1,x_2,\dots,x_\ell)\right)]\right|^2 \\
&\geqslant \left|\mathop{\mathbf{E}}_{x_1,x_2,\dots,x_\ell}[\mathsf{e}\left(P'(x_1,x_2,\dots,x_\ell)\right)]\right|^2 = \left|\mathop{\mathbf{E}}_{x_1,x_2,\dots,x_\ell}[\mathsf{e}\left(P(x_1,x_2,\dots,x_\ell)\right)]\right|^2.
\end{aligned}$$

□

□

The above proof also implies the following proposition just by omitting the application of Claim 4.6.

Proposition 4.9. *Let $L_1,\dots,L_m$ be linear forms on ℓ variables and let $\mathcal{B} = (P_1,\dots,P_C)$ be an ε-uniform polynomial factor of degree $d>0$ for some $\varepsilon\in(0,1]$ which is defined by only homogeneous polynomials. For every tuple Λ of integers $(\lambda_{i,j})_{i\in[C],j\in[m]}$, define*

$$P_\Lambda(X) = \sum_{i\in[C],j\in[m]} \lambda_{i,j}P_i(L_j(X)),$$

where $P_\Lambda : (\mathbb{F}^n)^\ell \to \mathbb{T}$. Moreover assume that for every $i\in[C], j\in[m]$, $L_j\in\mathcal{L}_{\deg(\lambda_{i,j}P_i)}$. Then, P_Λ is of degree $d=\max_{i,j}\deg(\lambda_{i,j}P_i)$ and $\|\mathsf{e}\left(P_\Lambda\right)\|_{U^d}<\varepsilon$.

4.3 The Gowers-Wolf Conjecture: Proof of Theorem 3.14

Theorem 3.14 (restated). *Let $\mathcal{L}=\{L_1,\dots,L_m\}$ be a system of linear forms. Assume that L_1^{d+1} is not in the linear span of $L_2^{d+1},\dots,L_m^{d+1}$. For every $\varepsilon>0$, there exists $\delta>0$ such that for any collection of functions $f_1,\dots,f_m:\mathbb{F}^n\to\mathbb{D}$ with $\|f_1\|_{U^{d+1}}\leqslant\delta$, we have*

$$\left|\mathop{\mathbf{E}}_{X\in(\mathbb{F}^n)^k}\left[\prod_{i=1}^m f_i(L_i(X))\right]\right| \leqslant \varepsilon. \tag{15}$$

We prove Theorem 3.14 in this section. Note that since L_1^{d+1} is not in the linear span of $L_2^{d+1},\dots,L_m^{d+1}$ we have that L_1 is linearly independent from each L_j for $j>1$. We claim that we may assume without loss of generality that $L_2,\dots,L_m$ are pairwise linearly independent as well, namely that $\{L_1,\dots,L_m\}$ has bounded Cauchy-Schwarz complexity. Assume that there are $j,\ell\in[m]\backslash\{1\}$ and a nonzero $c\in\mathbb{F}$ such that $L_\ell = cL_j$. Then we may define a new function

$f_j'(x) = f_j(x)f_\ell(cx)$ so that $f_j'(L_j(X)) = f_j(L_j(X))f_\ell(L_\ell(X))$ and remove the linear form L_ℓ and functions f_j, f_ℓ from the system. Now

$$\mathop{\mathbf{E}}_{X\in(\mathbb{F}^n)^k}\left[\prod_{i=1}^m f_i(L_i(X))\right] = \mathop{\mathbf{E}}_{X\in(\mathbb{F}^n)^k} f_j'(L_j(X))\cdot\left[\prod_{i\in[m]\setminus\{j,\ell\}} f_i(L_i(X))\right],$$

and thus it suffices to bound the right hand side of the above identity. We may repeatedly apply the above procedure in order to achieve a new system of pairwise linearly independent linear forms along with their corresponding functions while keeping L_1 and f_1 untouched.

Thus we may assume that $\{L_1,\dots,L_m\}$ is of finite Cauchy-Schwarz complexity s for some $s \leqslant m < \infty$. The case when $d \geqslant s$ follows from Lemma 3.12, thus we will consider the case when $d < s$. We will use Corollary 3.5 to write $f_i = g_i + h_i$ with

1. $g_i = \mathbf{E}[f_i|\mathcal{B}]$, where $\mathcal{B}$ is an r-regular polynomial factor of degree at most s and complexity $C \leqslant C_{\max}(p,s,\eta,\delta,r(\cdot))$ defined by only homogeneous polynomials, where r is a sufficiently fast growing growth function (to be determined later);

2. $\|h_i\|_{U^{s+1}} \leqslant \eta$.

We first show that by choosing a sufficiently small η we may replace f_i's in Equation (15) with g_i's.

Claim 4.10. *Choosing $\eta \leqslant \frac{\varepsilon}{2m}$ we have*

$$\left|\mathop{\mathbf{E}}_{X\in(\mathbb{F}^n)^k}\left[\prod_{i=1}^m f_i(L_i(X))\right] - \mathop{\mathbf{E}}_{X\in(\mathbb{F}^n)^k}\left[\prod_{i=1}^m g_i(L_i(X))\right]\right| \leqslant \frac{\varepsilon}{2}.$$

Proof. We have

$$\begin{aligned}\left|\mathop{\mathbf{E}}_X\left[\prod_{i=1}^m f_i(L_i(X))\right] - \mathop{\mathbf{E}}_X\left[\prod_{i=1}^m g_i(L_i(X))\right]\right| &= \left|\sum_{i=1}^m \mathop{\mathbf{E}}_X\left[h_i(L_i(X))\cdot\prod_{j=1}^{i-1} g_j(L_j(X))\cdot\prod_{j=i+1}^m f_j(L_j(X))\right]\right| \\ &\leqslant \sum_{i=1}^m\left|\mathop{\mathbf{E}}_X\left[h_i(L_i(X))\cdot\prod_{j=1}^{i-1} g_j(L_j(X))\cdot\prod_{j=i+1}^m f_j(L_j(X))\right]\right| \\ &\leqslant \sum_{i=1}^m \|h_i\|_{U^{s+1}} \leqslant m\cdot\eta \leqslant \frac{\varepsilon}{2},\end{aligned}$$

where the second inequality follows from Lemma 3.12 since the Cauchy-Schwarz complexity of $\mathcal{L}$ is s. □

Thus it is sufficient to bound $\left|\mathbf{E}_{X\in(\mathbb{F}^n)^k}\left[\prod_{i=1}^m g_i(L_i(X))\right]\right|$ by $\varepsilon/2$. For each i, $g_i = \mathbf{E}[f_i|\mathcal{B}]$ and thus

$$g_i(x) = \Gamma_i(P_1(x),\dots,P_C(x)),$$

where $P_1,\dots,P_C$ are the (non-classical) homogeneous polynomials of degree $\leqslant s$ defining $\mathcal{B}$ and $\Gamma_i : \mathbb{T}^C \to \mathbb{D}$ is a function. Let k_i denote the depth of the polynomial P_i so that by Lemma 2.6,

each P_i takes values in $\mathbb{U}_{k_i+1} = \frac{1}{p^{k_i+1}}\mathbb{Z}/\mathbb{Z}$. Moreover let $\Sigma := \mathbb{Z}_{p^{k_1+1}} \times \cdots \times \mathbb{Z}_{p^{k_C+1}}$. Using the Fourier expansion of Γ_i we have

$$g_i(x) = \sum_{\Lambda=(\lambda_1,\ldots,\lambda_C)\in\Sigma} \widehat{\Gamma}_i(\Lambda) \cdot \mathsf{e}\left(\sum_{j=1}^{C} \lambda_j P_j(x)\right), \tag{16}$$

where $\widehat{\Gamma}_i(\Lambda)$ is the Fourier coefficient of Γ_i corresponding to Λ. Let $P_\Lambda := \sum_{j=1}^{C} \lambda_j P_j(x)$ for the sake of brevity so that we may write

$$\mathop{\mathbf{E}}_{X}\left[\prod_{i=1}^{m} g_i(L_i(X))\right] = \sum_{\Lambda_1,\ldots,\Lambda_m\in\Sigma} \left(\prod_{i=1}^{m} \widehat{\Gamma}_i(\Lambda_i)\right) \cdot \mathop{\mathbf{E}}_{X}\left[\mathsf{e}\left(\sum_{i=1}^{m} P_{\Lambda_i}(L_i(X))\right)\right]. \tag{17}$$

We will show that for a sufficiently fast growing choice of the regularity function $r(\cdot)$ we may bound each term in Equation (17) by $\sigma := \frac{\varepsilon}{2|\Sigma|^m}$, thus concluding the proof by the triangle inequality. We will first show that the terms for which $\deg(P_{\Lambda_1}) \leqslant d$ can be made small.

Claim 4.11. *Let $\Lambda \in \Sigma$ be such that $\deg(P_\Lambda) \leqslant d$. For a sufficiently fast growing choice of $r(\cdot)$ and choice of $\delta \leqslant \frac{\sigma}{2}$,*

$$\left|\widehat{\Gamma}_1(\Lambda)\right| < \sigma.$$

Proof. It follows from Equation (16) that

$$\widehat{\Gamma}_1(\Lambda) = \mathop{\mathbf{E}}_{x}\left[g_1(x)\mathsf{e}\left(-P_\Lambda(x)\right)\right] - \sum_{\Lambda'\in\Sigma\setminus\{\Lambda\}} \widehat{\Gamma}_1(\Lambda') \cdot \mathop{\mathbf{E}}_{x}\left[\mathsf{e}\left(P_{\Lambda'}(x) - P_\Lambda(x)\right)\right].$$

Note that $\|f_1\|_{U^{d+1}} \leqslant \delta$ and thus

$$\left|\mathop{\mathbf{E}}_{x}\left[g_1(x)\mathsf{e}\left(-P_\Lambda(x)\right)\right]\right| = \left|\mathop{\mathbf{E}}_{x}\left[f_1(x)\mathsf{e}\left(-P_\Lambda(x)\right)\right]\right| \leqslant \|f_1\mathsf{e}\left(-P_\Lambda\right)\|_{U^{d+1}} = \|f_1\|_{U^{d+1}} \leqslant \delta \leqslant \sigma/2,$$

where we used of the fact that $g_1 = \mathbf{E}[f_1|\mathcal{B}]$ and the fact that Gowers norms are increasing in d. Finally the terms of the form $\widehat{\Gamma}_1(\Lambda') \cdot \mathbf{E}_x[\mathsf{e}\left(P_{\Lambda'}(x) - P_\Lambda(x)\right)]$ with $\Gamma' \neq \Gamma$ can be made arbitrarily small by choosing a sufficiently fast growing $r(\cdot)$ due to Remark 2.18, since $P_{\Lambda'} - P_\Lambda = P_{\Lambda'-\Lambda}$ is a nonzero linear combination of the polynomials defining the r-regular factor $\mathcal{B}$. □

The above claim allows us to bound the terms from Equation (17) corresponding to tuples $(\Lambda_1,\ldots,\Lambda_m) \in \Sigma^m$ with $\deg(P_{\Lambda_1}) \leqslant d$. This is because for such terms $|\widehat{\Gamma}_1(\Lambda_1)| < \sigma$ by the above claim, and $|\widehat{\Gamma}_i(\Lambda_i)| \leqslant 1$ since f_i's take values in $\mathbb{D}$. It remains to bound the terms for which $\deg(P_{\Lambda_1}) > d$. We will need the following claim.

Lemma 4.12. *Assume that L_1^{d+1} is not in the linear span of $L_2^{d+1},\ldots,L_m^{d+1}$, and let $(\Lambda_1,\ldots,\Lambda_m) \in \Sigma^m$ be such that $\deg(P_{\Lambda_1}) \geqslant d+1$. Then*

$$\sum_{i=1}^{m} P_{\Lambda_i}(L_i(X)) \not\equiv 0.$$

This combined with Theorem 3.7 implies that for a sufficiently fast growing choice of $r(\cdot)$, $\mathbf{E}_X\left[\mathsf{e}\left(\sum_{i=1}^{m} P_{\Lambda_i}(L_i(X))\right)\right] < \sigma$ which completes the proof of Theorem 3.14. Thus, we are left with proving Lemma 4.12.

Proof of Lemma 4.12: Assume to the contrary that $\sum_{i=1}^{m} P_{\Lambda_i}(L_i(X)) \equiv 0$. Denoting the coordinates of Λ_i by $(\lambda_{i,1}, \ldots, \lambda_{i,C}) \in \Sigma^C$ we have

$$\sum_{i=1}^{m} P_{\Lambda_i}(L_i(X)) = \sum_{i\in[m], j\in[C]} \lambda_{i,j} P_j(L_i(X)) \equiv 0.$$

Since the polynomial factor defined by $P_1, \ldots, P_C$ is r-regular with a sufficiently fast growing growth function $r(\cdot)$, Theorem 3.7 implies that for every $j \in [C]$ we must have that

$$\sum_{i=1}^{m} \lambda_{i,j} P_j(L_i(X)) \equiv 0. \tag{18}$$

Since $\deg(P_{\Lambda_1}) > d$, there must exist $j \in [C]$ such that $\lambda_{1,j} \neq 0$ and $\deg(\lambda_{1,j} P_j) > d$. Let $j^* \in [C]$ be such that $\deg(\lambda_{1,j^*} P_{j^*})$ is maximized, and let $d^* := \deg(P_{j^*})$ (note that $\deg(\lambda_{1,j^*} P_{j^*}) \leqslant d^*$). We will first prove that replacing P_{j^*} with a classical homogeneous polynomial Q of the same degree, Equation (18) for $j = j^*$ would still hold.

Claim 4.13. *Let $Q : \mathbb{F}^n \to \mathbb{T}$ be a classical homogeneous polynomial with $\deg(Q) = d^*$. Then*

$$\sum_{i=1}^{m} \lambda_{i,j^*} Q(L_i(X)) \equiv 0.$$

Proof. Assume to the contrary that $\sum_{i=1}^{m} \lambda_{i,j^*} Q(L_i(X)) \not\equiv 0$. By Claim 4.5 for degree d^* and linear forms $L_1, \ldots, L_m$, we can find a set of coefficients $\{a_{i,t} \in \mathbb{Z}\}_{i\in[\ell], t\in[m']}$, $\{c_{i,t} \in \mathbb{F}\backslash\{0\}\}_{i\in[\ell], t\in[m']}$ and linear forms $\{M_t\}_{t\in[m']}$ with $|M_t| \leqslant d^*$ such that the first nonzero entry of every M_t is equal to 1, such that

$$\sum_{i=1}^{m} \lambda_{i,j^*} Q(L_i(X)) = \sum_{i\in[\ell], t\in[m']} a_{i,t} \lambda_{i,j^*} Q(c_{i,t} M_t(X)) = \sum_{t=1}^{m'} \alpha_t Q(M_t(X)) \not\equiv 0, \tag{19}$$

where $\alpha_t = \sum_{i=1}^{\ell} a_{i,t} \lambda_{i,j^*} |c_{i,t}|^{d^*}$. Here the second equality follows from Q being classical and homogeneous. Furthermore,

$$\sum_{i=1}^{m} \lambda_{i,j^*} P_{j^*}(L_i(X)) = \sum_{i\in[\ell], t\in[m']} a_{i,t} \lambda_{i,j^*} P_{j^*}(c_{i,t} M_t(X)) = \sum_{t=1}^{m'} \beta_t P_{j^*}(M_t(X)), \tag{20}$$

where $\beta_t = \sum_{i=1}^{\ell} a_{i,t} \lambda_{i,j^*} \sigma_{i,t}$, $\{\sigma_{i,t}\}_{i\in[\ell], t\in[m']}$ are integers whose existence follows from the homogeneity of P_{j^*}. Moreover, by Remark 3.2 we know that $\sigma_{i,t} \equiv |c_{i,t}|^{d^*} \mod p$, and hence

$$\alpha_t \equiv \beta_t \mod p, \quad \forall t \in [m']. \tag{21}$$

Now notice that Equation (19) implies that there exists some $t \in [m']$ for which $\alpha_t Q \neq 0$, which, since Q is classical, is equivalent to $\alpha_t \not\equiv 0 \mod p$. Hence also $\beta_t \not\equiv 0 \mod p$. Let $T = \{t \in [m'] : \beta_t \not\equiv 0 \mod p\}$, which we just verified is nonempty. Then for $t \in T$, $\deg(\beta_t P_{j^*}) = \deg(P_{j^*}) = d^*$; and for $t \in [m] \setminus T$, $\deg(\beta_t P_{j^*}) \leqslant d^* - (p-1) < d^*$.

We can now decompose

$$\sum_{i=1}^{m} \lambda_{i,j^*} P_{j^*}(L_i(X)) = \sum_{t \in T} \beta_t P_{j^*}(M_t(X)) + \sum_{t \in [m] \setminus T} \beta_t P_{j^*}(M_t(X)). \tag{22}$$

By Proposition 4.9, the first sum in Equation (22) is a nonzero polynomial of degree d^*, and by our previous argument, the sum for $t \notin T$ is a polynomial of degree less than d^*. Hence, we get that $\sum_{i=1}^{m} \lambda_{i,j^*} P_{j^*}(L_i(X))$ is a nonzero polynomial of degree d^*, which is a contradiction to our assumption. □

We have proved that for every choice of a degree-d^* classical homogeneous polynomial Q, $\sum_{i=1}^{m} \lambda_{i,j^*} Q(L_i(X)) \equiv 0$. Now choosing the polynomial $Q(x(1), \ldots, x(n)) = x(1) \cdots x(d^*)$ and looking at the coefficients of the monomials of degree d^*, we have

$$\sum_{i=1}^{m} \lambda_{i,j^*} L_i^{d^*} \equiv 0.$$

Recalling that $\lambda_{1,j^*} \neq 0$ and that $d^* > d$, this means that L_1^{d+1} can be written as a linear combination of $L_2^{d+1}, \ldots, L_m^{d+1}$, a contradiction. □

References

[BFH+13] Arnab Bhattacharyya, Eldar Fischer, Hamed Hatami, Pooya Hatami, and Shachar Lovett, *Every locally characterized affine-invariant property is testable*, Proceedings of the 45th annual ACM symposium on Symposium on theory of computing (New York, NY, USA), STOC '13, ACM, 2013, pp. 429–436. 2, 3, 8, 10, 11, 16, 19

[BFL13] Arnab Bhattacharyya, Eldar Fischer, and Shachar Lovett, *Testing low complexity affine-invariant properties*, Proc. 24th ACM-SIAM Symposium on Discrete Algorithms, 2013, pp. 1337–1355. 8

[BTZ10] Vitaly Bergelson, Terence Tao, and Tamar Ziegler, *An inverse theorem for the uniformity seminorms associated with the action of* $\mathbb{F}_p^\infty$, Geom. Funct. Anal. **19** (2010), no. 6, 1539–1596. MR 2594614 (2011b:37009) 2

[Gow01] W. T. Gowers, *A new proof of Szemerédi's theorem*, Geom. Funct. Anal. **11** (2001), no. 3, 465–588. MR 1844079 (2002k:11014) 1

[Gow10] ———, *Decompositions, approximate structure, transference, and the Hahn-Banach theorem*, Bull. Lond. Math. Soc. **42** (2010), no. 4, 573–606. MR 2669681 (2011k:11015) 8

[Gre07] Ben Green, *Montréal notes on quadratic Fourier analysis*, Additive combinatorics, CRM Proc. Lecture Notes, vol. 43, Amer. Math. Soc., Providence, RI, 2007, pp. 69–102. MR 2359469 (2008m:11047) 8

[GT09] Ben Green and Terence Tao, *The distribution of polynomials over finite fields, with applications to the Gowers norms*, Contrib. Discrete Math. **4** (2009), no. 2, 1–36. MR 2592422 (2011d:11022) 2, 8, 10

[GT10a] ———, *An arithmetic regularity lemma, an associated counting lemma, and applications*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 261–334. MR 2815606 2, 3

[GT10b] Benjamin Green and Terence Tao, *Linear equations in primes*, Ann. of Math. (2) **171** (2010), no. 3, 1753–1850. MR 2680398 (2011j:11177) 12, 13

[GTZ12] Ben Green, Terence Tao, and Tamar Ziegler, *An inverse theorem for the Gowers $U^{s+1}[N]$-norm*, Ann. of Math. (2) **176** (2012), no. 2, 1231–1372. MR 2950773 2

[GW10] W. T. Gowers and J. Wolf, *The true complexity of a system of linear equations*, Proc. Lond. Math. Soc. (3) **100** (2010), no. 1, 155–176. MR 2578471 (2011a:11019) 3, 13

[GW11] ———, *Linear forms and higher-degree uniformity for functions on $\mathbb{F}_p^n$*, Geom. Funct. Anal. **21** (2011), no. 1, 36–69. MR 2773103 (2012f:11024) 3, 13

[HL11] Hamed Hatami and Shachar Lovett, *Higher-order Fourier analysis of $\mathbb{F}_p^n$ and the complexity of systems of linear forms*, Geom. Funct. Anal. **21** (2011), no. 6, 1331–1357. MR 2860190 (2012m:11016) 2, 10, 13

[KL08] Tali Kaufman and Shachar Lovett, *Worst case to average case reductions for polynomials*, Foundations of Computer Science, IEEE Annual Symposium on **0** (2008), 166–175. 2

[Sam07] Alex Samorodnitsky, *Low-degree tests at large distances*, STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing, ACM, New York, 2007, pp. 506–515. MR MR2402476 (2009f:68077) 2

[Sze12] Balazs Szegedy, *On higher order fourier analysis*, http://arxiv.org/1203.22609. 2

[Tao07] Terence Tao, *Structure and randomness in combinatorics*, Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on, IEEE, 2007, pp. 3–15. 8

[TZ10] Terence Tao and Tamar Ziegler, *The inverse conjecture for the Gowers norm over finite fields via the correspondence principle*, Anal. PDE **3** (2010), no. 1, 1–20. MR 2663409 (2011j:11018) 2

[TZ12] ———, *The inverse conjecture for the Gowers norm over finite fields in low characteristic*, Ann. Comb. **16** (2012), no. 1, 121–188. MR 2948765 2, 4, 5, 6, 8, 10